NOAA R&DHPCS ESRL - Boulder

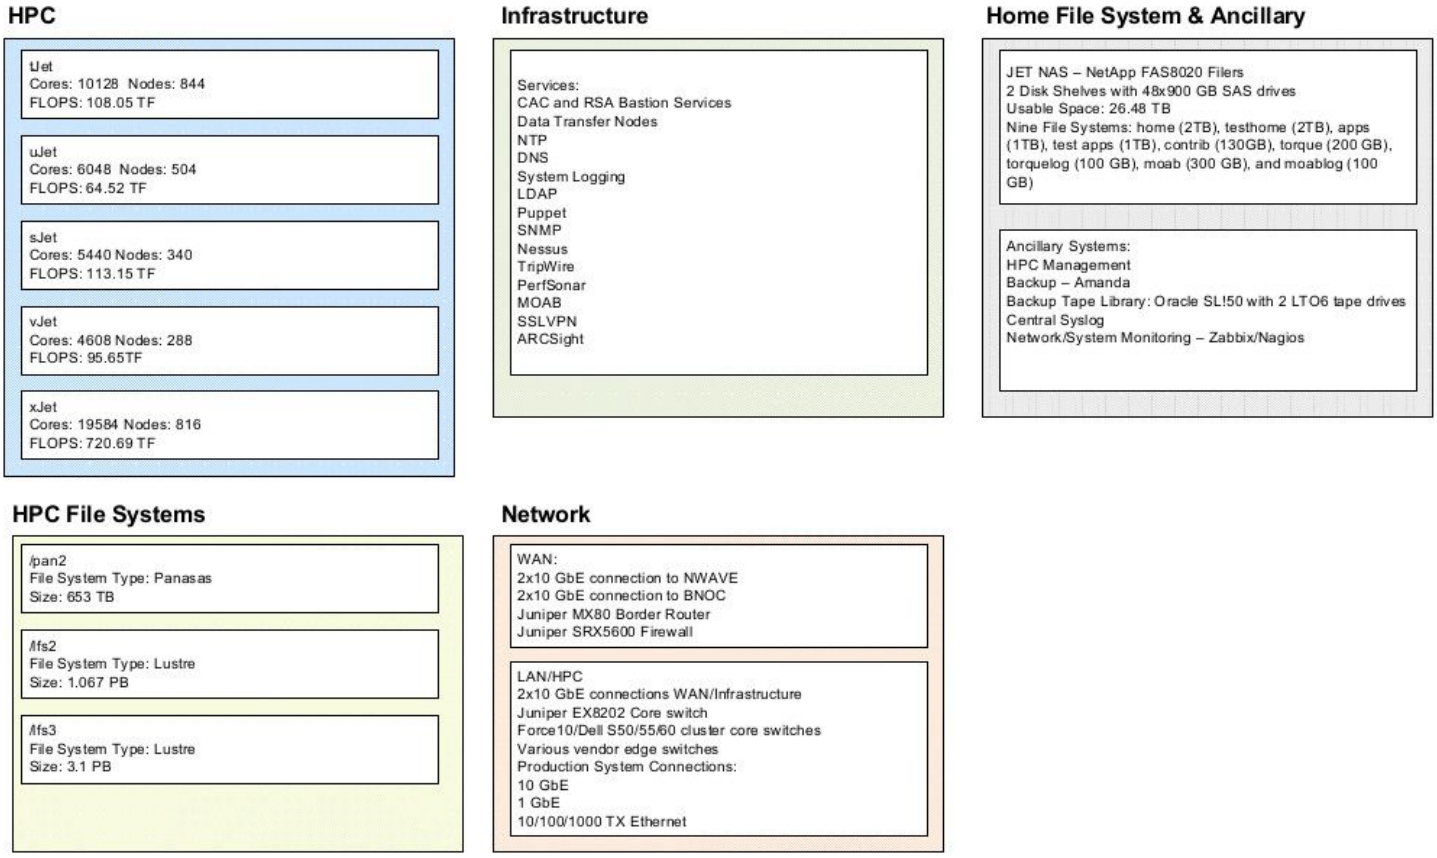

Detailed HPCS Architecture Boulder

16 January 2017

Boulder Facility Organization and Subsystems Overview

The Boulder HPC systems are divided amongst two rooms in the David Skaggs Research Center (DSRC) in Boulder, CO. Room 2B407 is primarily for low power density air cooled equipment, such as parallel file systems, Front Ends/Login nodes, Batch nodes, Infrastructure servers, and Home File System. There is also a 288 node compute cluster located in 2B407. However, this cluster requires a special air curtain for cold air isolation.

Room GA405 is for high density air cooled equipment such as compute racks. Each compute rack is comprised of 60-70 nodes, IB switches and Ethernet switches. Four of the five compute clusters are located on the GA405 room. Both rooms are connected via 18 active, 150 meter, QDR IB links. So all compute systems reside on the same IB fabric.

Jet is comprised of 5 separate clusters with a total core count of 45,808 cores. A single job cannot span multiple clusters. All clusters can be accessed through the Batch system from the common Front-ends/Login nodes.

System	Intel Processor	Install Date	Clock (GHz)	Nodes	Total Cores	Memory Per Node (GB)	Memory Per Core (GB)
tJet	Westmere	8/2010	2.66	758	9096	24	2
uJet	Westmere	11/2011	2.66	590	7080	24	2
sJet	Sandy Bridge	8/2012	2.6	340	5440	32	2
vJet	Ivy Bridge	8/2014	2.6	288	4608	64	4
xJet	Haswell	9/2015 8/2016*	2.3	816	19584	64	2.67

* xJet expansion

Jet has approximately 4.82PB of scratch storage. This is across 3 parallel file-systems (2 DDN Lustre and 1 Panasas). Combined peak performance of all three scratch file systems is ~46GB/s. All clusters reside on the same High Speed Interconnect/Infiniband fabric, making the parallel file systems globally accessible from all nodes. There is also a Home File System (HFS) which is globally accessible from all nodes over the Ethernet network.

The Home File System (HFS) and other NFS file systems are being served by a NetApp FAS8120. The HFS is globally accessible from all nodes over the Ethernet network.

The HFS and HPC infrastructure servers are backed up on a daily and weekly basis to a combination of spinning disk and tape. The Open Source software Amanda is used for all backups.

There is a tape storage archive located at the NESCC site in Fairmont, WV. Data can be transferred to and from this archive from Jet's Front Ends. The retention period for the archive is 1-5 years or permanent. All projects with a compute allocation are allowed to store data on the archive.

Data transfer services are provided by a combination of 10GbE connected Data Transfer Nodes (DTN's) and the 10GbE connected Front Ends. The DTN's are used when a data transfer is initiated from outside of Jet. The Jet FE's are used when a data transfer is initiated from Jet.

Jet's Batch system is comprised of three separate pieces of software. Torque is the resource manager and is responsible for job submission, node health status and job launching. Moab is the scheduler and is responsible for scheduling jobs based on priority and managing dedicated system resources with reservations. MAM in the allocation manager and is responsible for managing the monthly allocation of core hours for each project. MAM and Moab are also used for the daily and monthly reporting of system utilization.

Diagrams

NOAA/Boulder

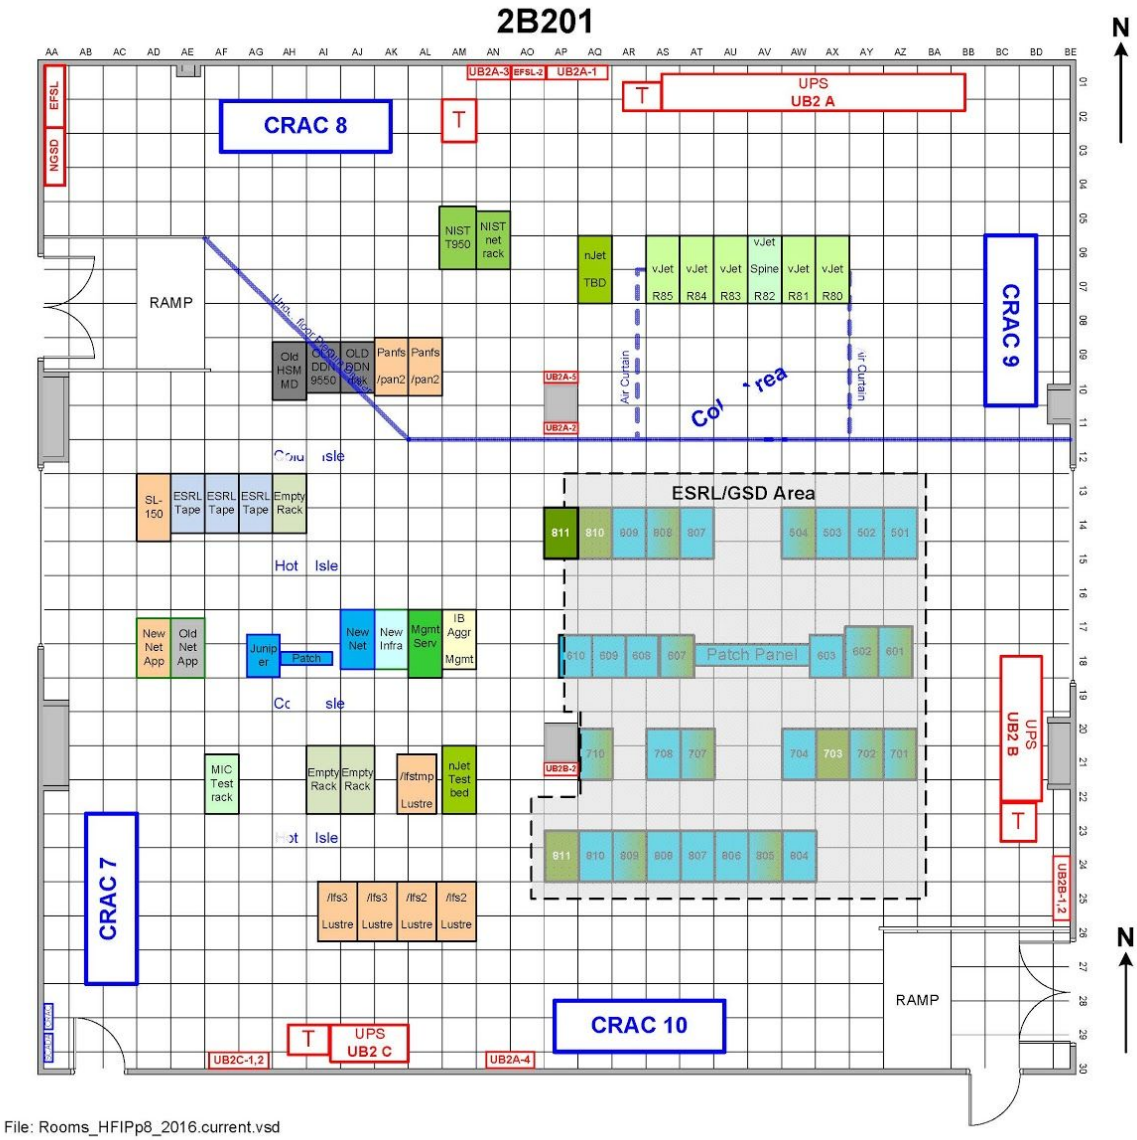

2B201
~08/01/2016

HFIP p8

File: Rooms_HFIPp8_2016.current.vsd

GA405
~08/01/2016

HFIP p8

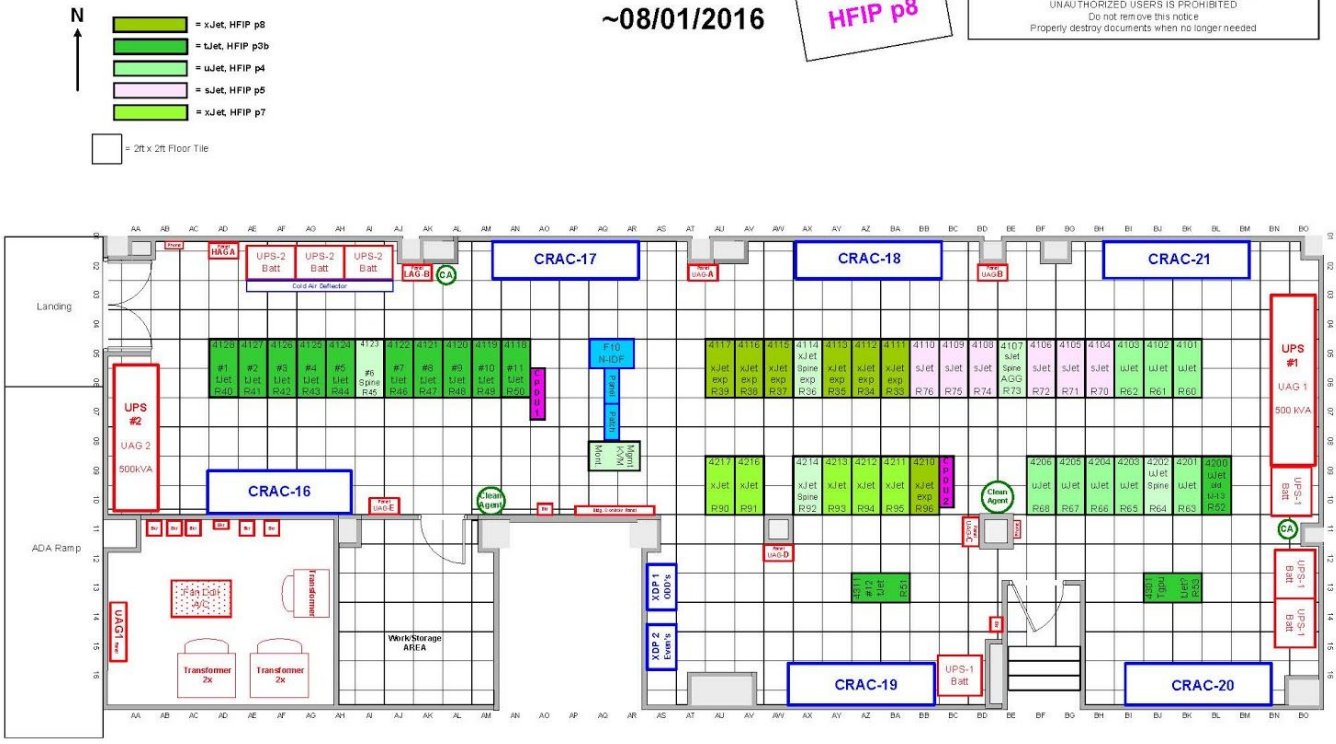

Boulder – External Connectivity

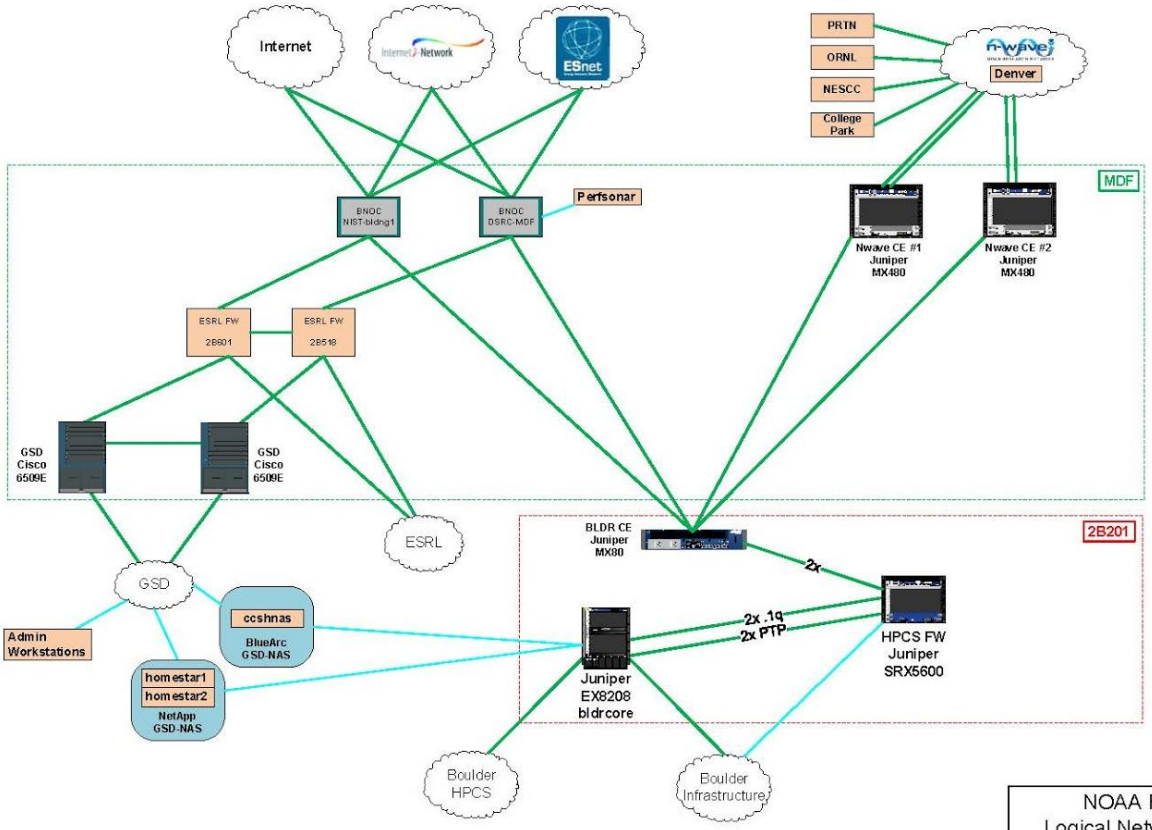

Legend

- Ten Gigabit Ethernet
- Gigabit Ethernet
- 10/100/1000TX Ethernet

NOAA R&DHPCS Logical Network Diagrams	
Date: 10/31/13	Version: 5.3
Drawn For: National Oceanic and Atmospheric Administration	
	Pg: 3

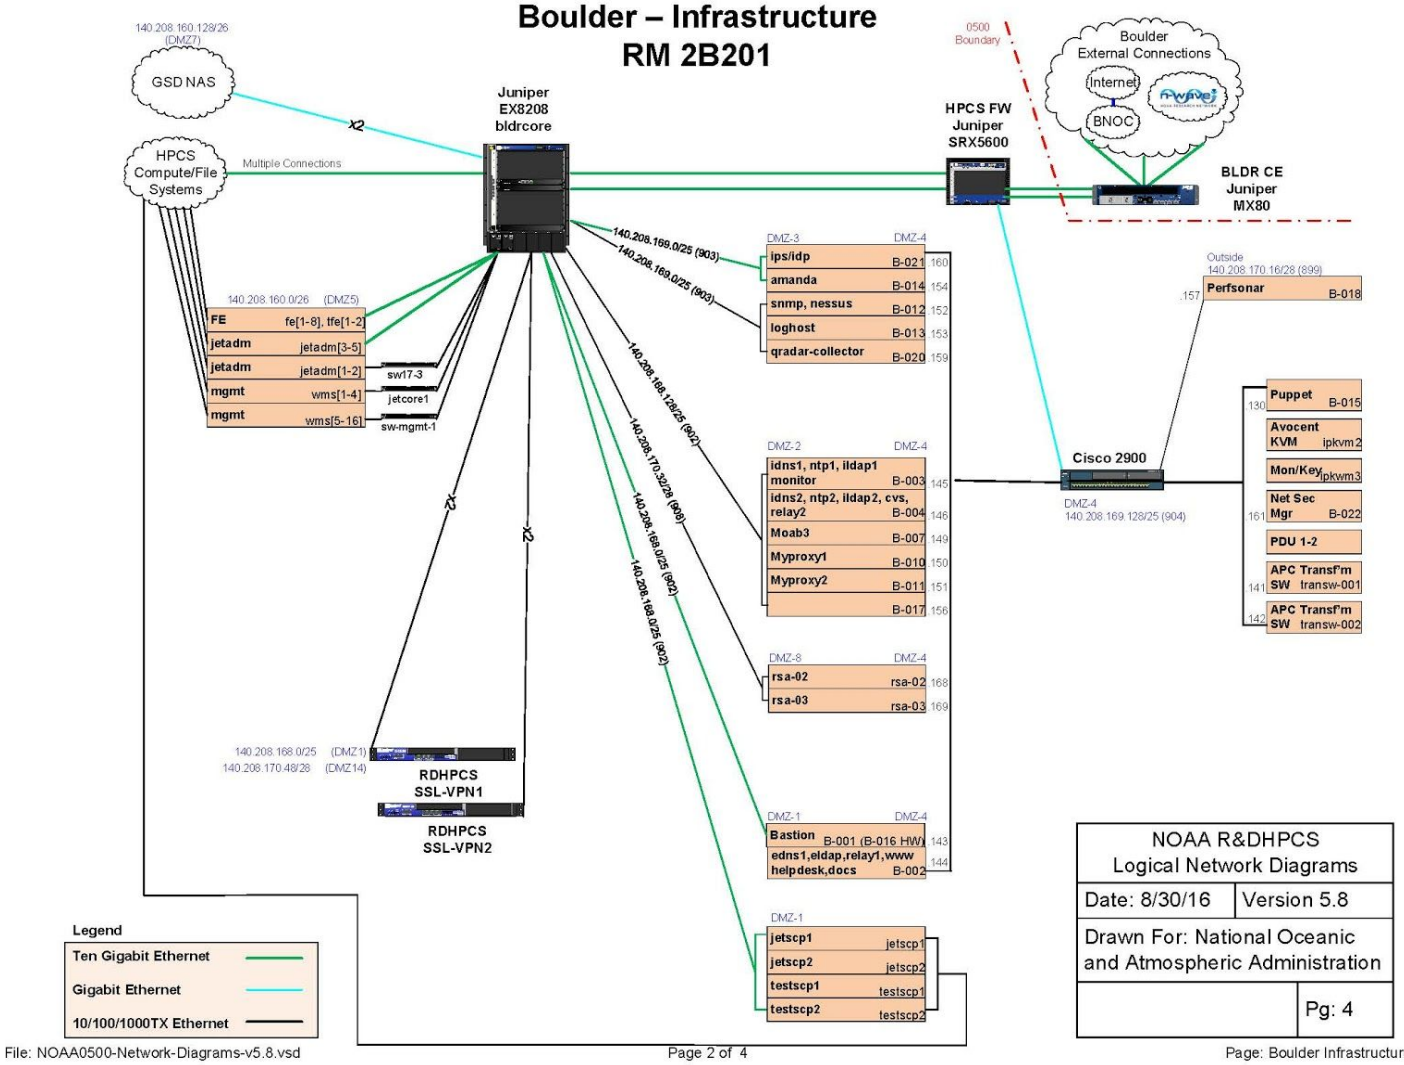

NOAA R&DHPCS Logical Network Diagrams	
Date: 8/30/16	Version 5.8
Drawn For: National Oceanic and Atmospheric Administration	
	Pg: 4

Boulder – HPCS
RM GA405

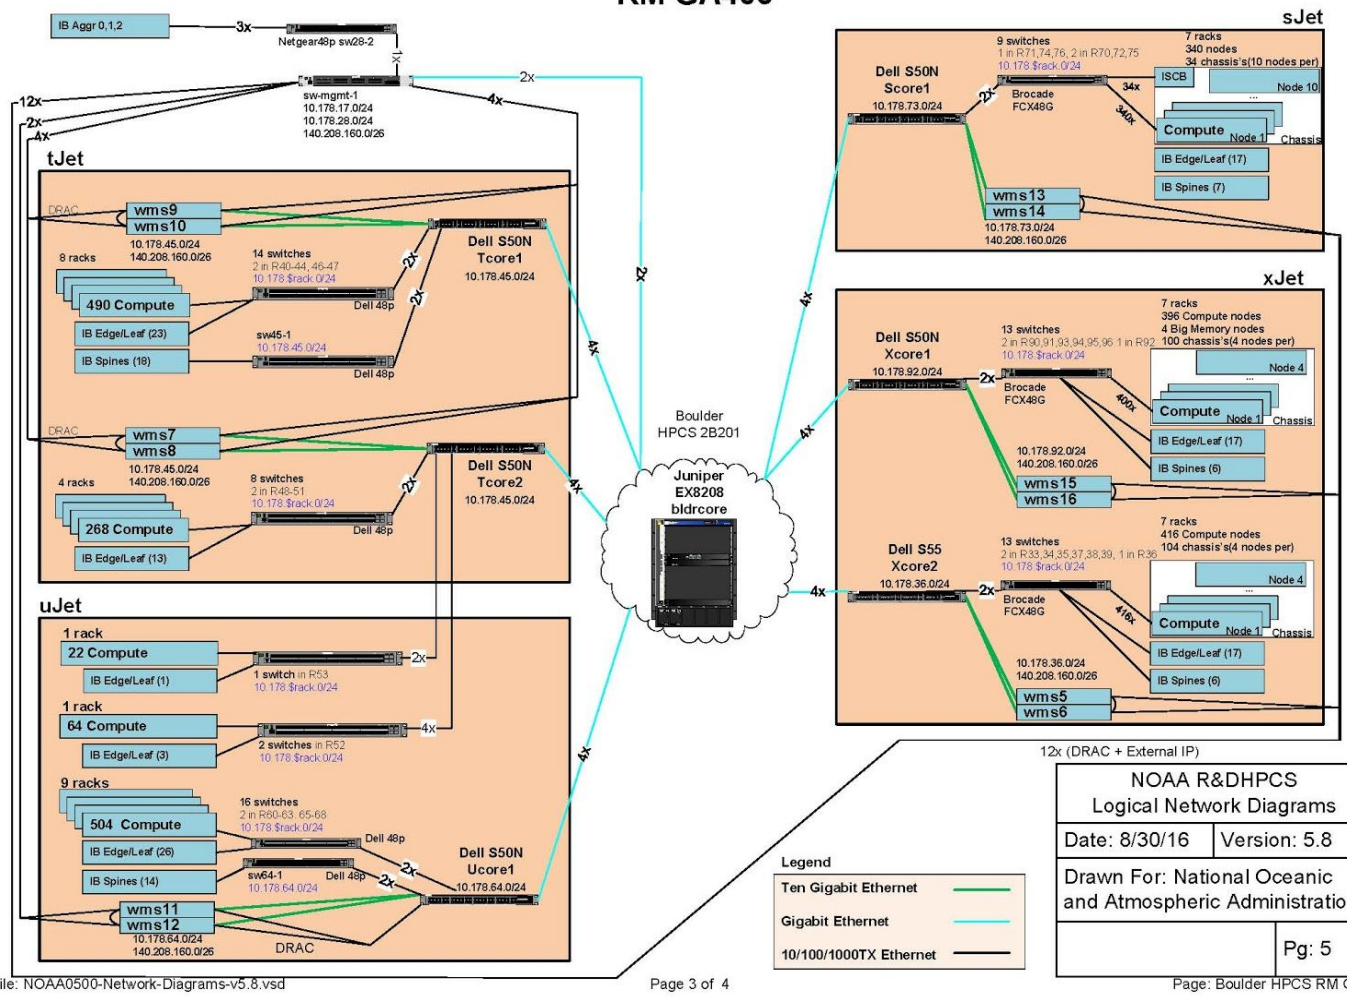

NOAA R&DHPCS Logical Network Diagrams	
Date: 8/30/16	Version: 5.8
Drawn For: National Oceanic and Atmospheric Administration	
	Pg: 5

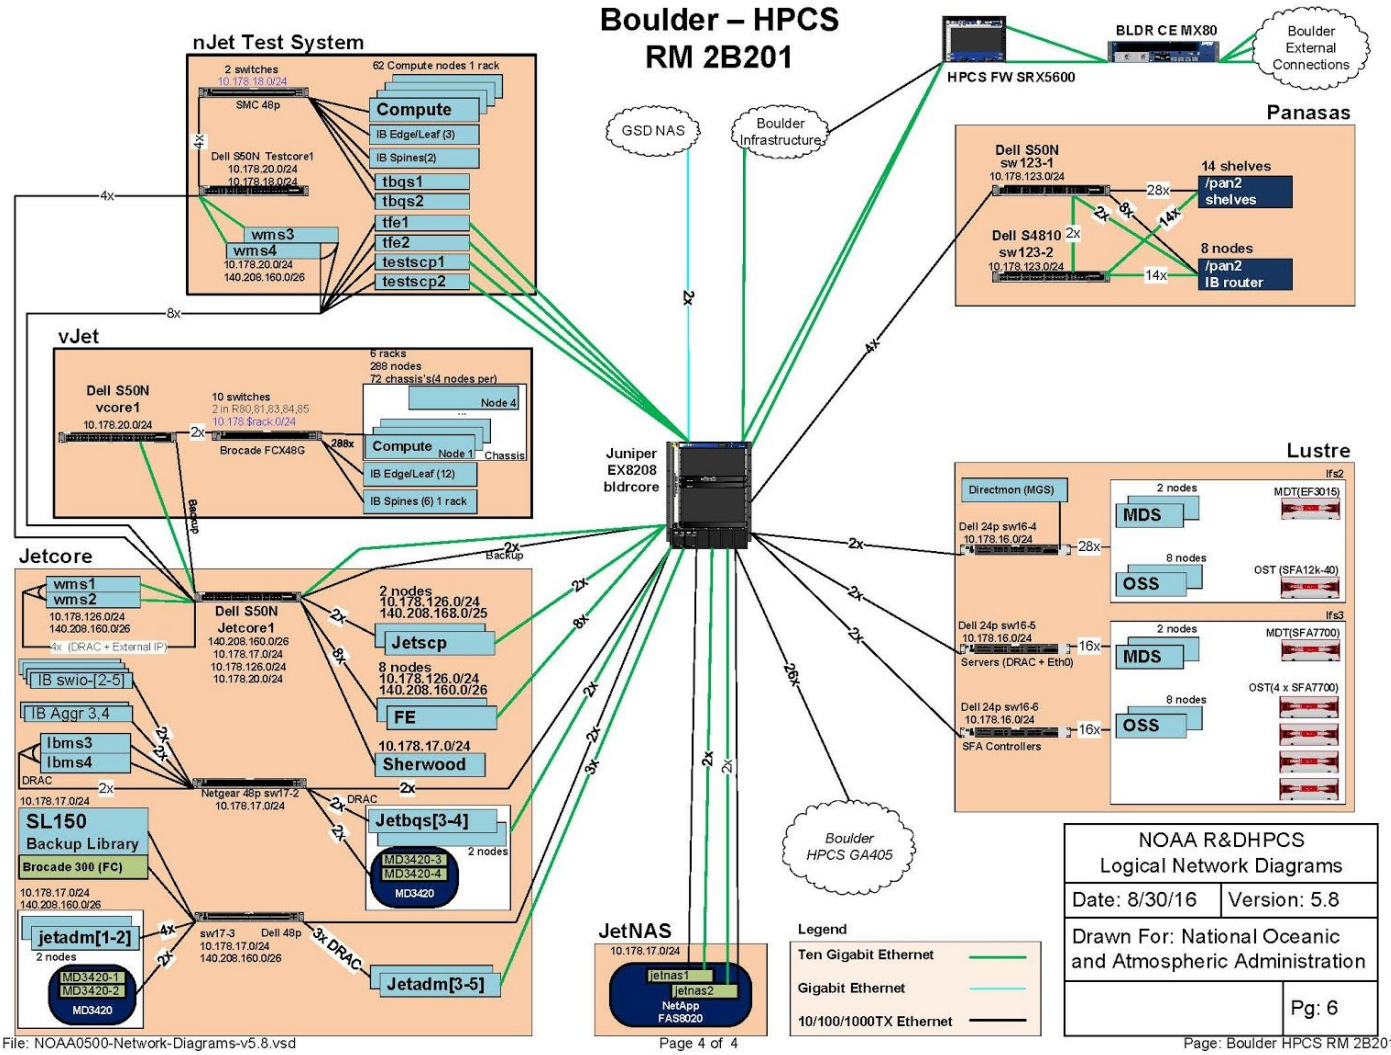
Boulder – HPCS
RM 2B201
nJet Test System
2 switches
10.178.18.0/24
SMC 48p
62 Compute nodes 1 rack
Compute
IB Edge/Leaf (3)
IB Spines(2)
Dell S50N Testcore1
10.178.20.0/24
10.178.18.0/24
tbqs1
tbqs2
tfe1
tfe2
testscp1
testscp2
wms3
wms4
10.178.20.0/24
140.208.160.0/26
vJet
Dell S50N
vcore1
10.178.20.0/24
10 switches
2 in R80,81,83,84,85
10.178.$rack.0/24
Brocade FCX48G
6 racks
288 nodes
72 chassis's(4 nodes per)
Node 4
Node 1
Chassis
Compute
IB Edge/Leaf (12)
IB Spines (6) 1 rack
Backup
Jetcore
wms1
wms2
10.178.126.0/24
140.208.160.0/26
4x (DRAC + External IP)
Dell S50N
Jetcore1
140.208.160.0/26
10.178.17.0/24
10.178.126.0/24
10.178.20.0/24
IB swio-[2-5]
IB Aggr 3,4
lbms3
lbms4
DRAC
10.178.17.0/24
Netgear 48p sw17-2
10.178.17.0/24
2 nodes
10.178.126.0/24
140.208.168.0/25
Jetscp
8 nodes
10.178.126.0/24
140.208.160.0/26
FE
10.178.17.0/24
Sherwood
Jetbqs[3-4]
2 nodes
MD3420-3
MD3420-4
MD3420
SL150
Backup Library
Brocade 300 (FC)
10.178.17.0/24
140.208.160.0/26
jetadm[1-2]
2 nodes
MD3420-1
MD3420-2
MD3420
sw17-3 Dell 48p
10.178.17.0/24
140.208.160.0/26
Jetadm[3-5]
GSD NAS
Boulder
Infrastructure
Juniper
EX8208
bldrcore
HPCS FW SRX5600
BLDR CE MX80
Boulder
External
Connections
Panasas
Dell S50N
sw123-1
10.178.123.0/24
14 shelves
/pan2
shelves
Dell S4810
sw123-2
10.178.123.0/24
8 nodes
/pan2
IB router
Lustre
Directmon (MGS)
Dell 24p sw16-4
10.178.16.0/24
2 nodes
MDS
MDT(EF3015)
8 nodes
OSS
OST (SFA12k-40)
lfs2
Dell 24p sw16-5
10.178.16.0/24
Servers (DRAC + Eth0)
Dell 24p sw16-6
10.178.16.0/24
SFA Controllers
MDT(SFA7700)
OST(4 x SFA7700)
lfs3
Boulder
HPCS GA405
JetNAS
10.178.17.0/24
jetnas1
jetnas2
NetApp
FAS8020
Legend
Ten Gigabit Ethernet
Gigabit Ethernet
10/100/1000TX Ethernet
NOAA R&DHPCS
Logical Network Diagrams
Date: 8/30/16
Version: 5.8
Drawn For: National Oceanic and Atmospheric Administration
Pg: 6

NOAA/Boulder

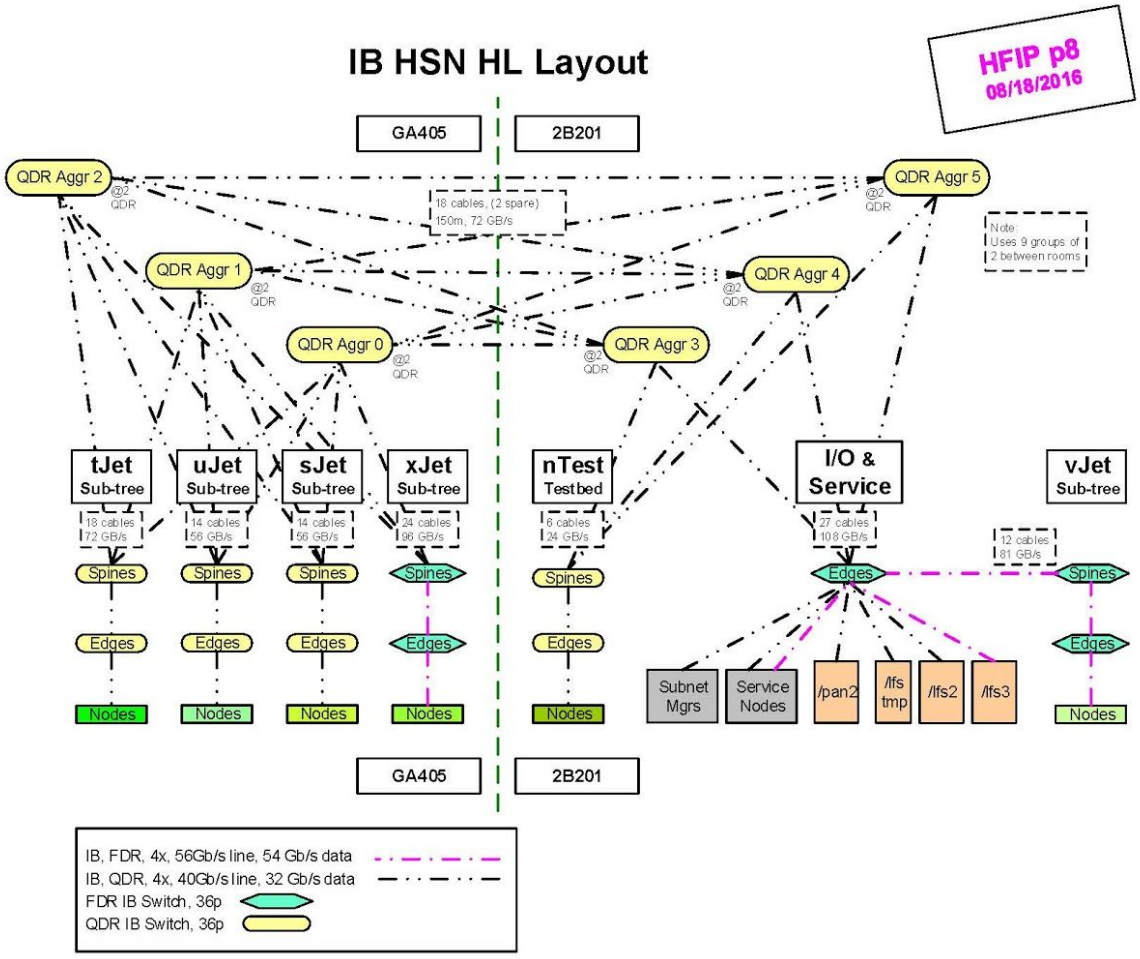

File: IB_Jet_HFIPp8_2016.08.18.vsd

Page: IB-H

NOAA/Boulder

IB HSN Layout - 2B201

HFIP p8
08/18/2016

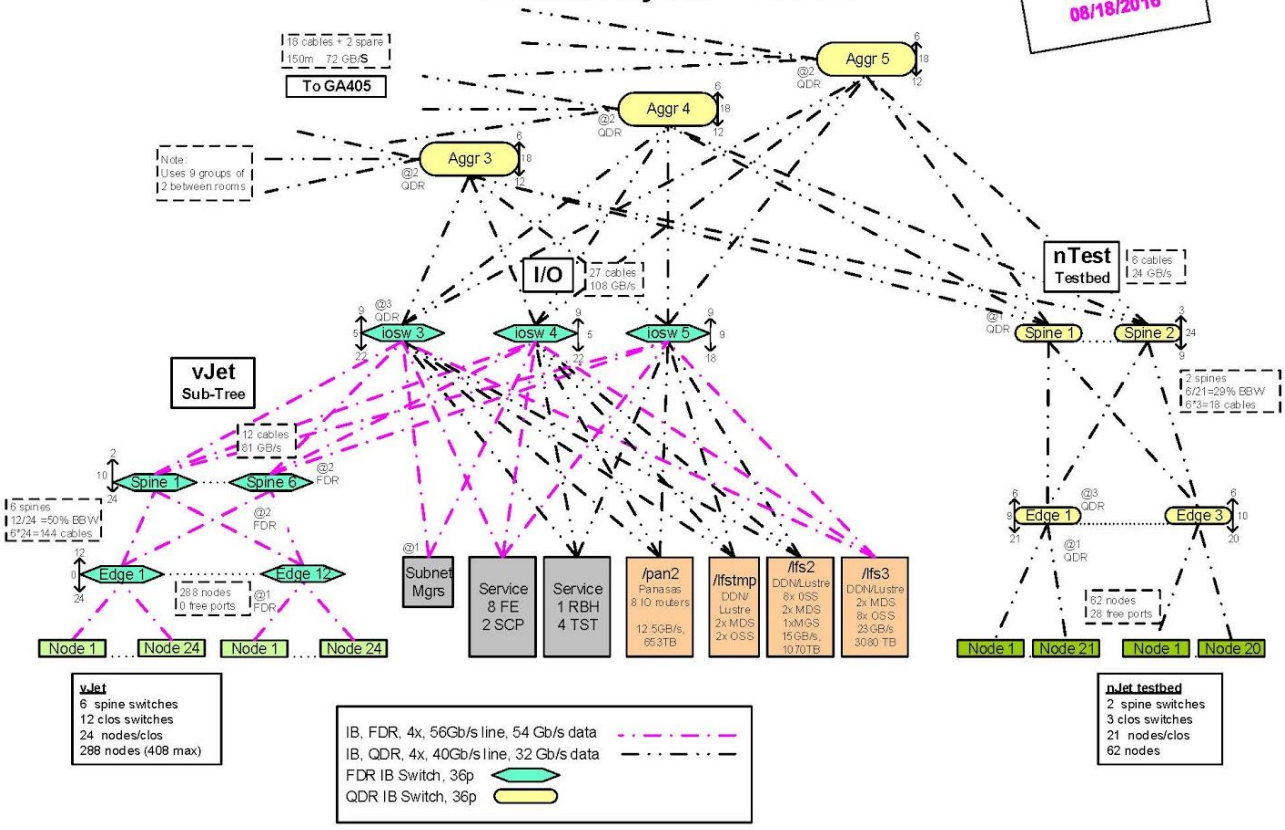

NOAA/Boulder

IB HSN Layout - GA405

HFIP p8
08/18/2016